

The weak Bruhat order for random walks on Coxeter groups

Graham White

November 15, 2016

Abstract

We show that for the simple random walk on a Coxeter group generated by the Coxeter generators and identity, the likelihoods of being at any pair of states respect the weak Bruhat order. That is, after any number of steps, the most likely element is the identity, probabilities decrease along any geodesic from the identity, and the least likely element is the longest element, if the group is finite. The result remains true when different generators have different probabilities, so long as the identity is at least as likely as any other.

1 Introduction

If a random walk is started in a known state and run for several steps, one may examine the probabilities that it is now in each possible state. A likelihood order is a partial order on the state space, so that if one state is larger than another, then the random walk is always more likely to be in the former state than the latter, after any number of steps. The main result of this paper is that for any Coxeter system (W, S) , the weak Bruhat order is a likelihood order for the simple random walk on W generated by $S \cup \{1_W\}$.

Theorem 1. *For any Coxeter system (W, S) , consider the simple random walk on W , starting at the identity, and at each step multiplying on the right by an element of S or by the identity, each with probability $\frac{1}{|S|+1}$. Then for any n , and any two states w and w' , if $w \leq_B w'$ then the probability that the random walk is at w after n steps is at least the probability that it is at w' .*

In type A, this result describes a partial order for the (appropriately lazy) adjacent transposition walk on the symmetric group. For analysis of the mixing time of this walk, see Section 4 of [6]. The adjacent transposition walk is a special case of the interchange process [8].

Likelihood orders can describe the most and least likely states of a random walk. In particular, if a random walk has uniform stationary distribution, then the separation distance from the stationary distribution depends only on the probability of being at the least likely state. Thus, knowing which state is the least likely, together with a lower bound on the probability of being at that state, produces an upper bound on the separation distance mixing time. Upper bounds on separation distance give upper bounds on total variation distance, so upper bounds on total variation mixing times follow.

Lower bounds on mixing times are obtained by analysing a set of unlikely states. Knowledge of which states are the least likely via a likelihood order can inform the choice of such a set. For instance, see Section 4 of [1].

Some results regarding likelihood orders for various random walks on the symmetric group S_n are given in [2] and [3]. These likelihood orders are shown to hold after enough steps (for example, after $O(n^2)$ steps). Likelihood orders have also been considered by Diaconis and Isaacs in [5], where they prove that for any symmetric random walk on a group, after any even number of steps the most likely state is the initial one. They also give likelihood orders for several random walks on the cycle.

2 Preliminaries

In this section, some necessary background results are recalled. Section 2.1 defines Cayley graphs and a useful family of their symmetries, and Section 2.2 discusses the sets of vertices in the Cayley graph which are closer to one end of a given edge than to the other. Section 2.3 defines the weak Bruhat order.

A Coxeter system (W, S) is a group W together with a presentation of a certain form.

Definition 2. A Coxeter presentation is a presentation of the form

$$\langle s_1, s_2, \dots, s_n \mid \{s_i^2\}_{i=1}^n, \{(s_i s_j)^{m_{ij}}\}_{i \neq j} \rangle$$

where each $m_{ij}, i \neq j$ is either a positive integer at least two, or ∞ , indicating the lack of that relation.

A good example of a Coxeter group is the symmetric group S_n , which has the Coxeter presentation

$$S_n = \left\langle s_1, s_2, \dots, s_{n-1} \left| \begin{array}{ll} s_i^2 & \text{for each } 1 \leq i \leq n-1 \\ (s_i s_{i+1})^3 & \text{for each } 1 \leq i \leq n-2 \\ s_i s_j s_i^{-1} s_j^{-1} & \text{if } |i-j| > 1 \end{array} \right. \right\rangle.$$

In this presentation, the generator s_i represents the transposition $(i \ i+1)$. This presentation has $m_{ij} = 3$ when $|i-j| = 1$ and $m_{ij} = 2$ for $|i-j| > 1$.

2.1 Cayley graphs

Given a group W and a generating set S , the Cayley graph $\Gamma(W, S)$ is defined as follows

Definition 3. *The graph $\Gamma(W, S)$ has a vertex for each element of W , and for each $w \in W$ and each $s \in S$, there is an edge from w to ws . It will often be convenient to label the edge (w, ws) by the generator s .*

In the present setting, groups will always be generated by elements of order two, so Cayley graphs will be undirected.

It will be necessary to have the following results regarding certain symmetries of Cayley graphs.

Definition 4. *Consider a Cayley graph $\Gamma(W, S)$. For any $x \in W$, let L_x be the left multiplication map on $\Gamma(W, S)$ which takes w to xw , for each $w \in W$.*

Proposition 5. *For any Cayley graph $\Gamma(W, S)$ and any $x \in W$, the map L_x is an automorphism of $\Gamma(W, S)$. Further, L_x preserves the edge labels of $\Gamma(W, S)$.*

Proof. The map $L_{x^{-1}}$ is the inverse of L_x , so L_w is a bijection. To check that L_x preserves edges of $\Gamma(W, S)$, observe that for each edge (w, ws) of $\Gamma(W, S)$, the image under L_x , (xw, xws) , is also an edge of $\Gamma(W, S)$, and that these two edges have the same label. \square

Random walks on the group W can be understood via the Cayley graph. In particular, if a random walk is defined by at each step multiplying by an element of S , then consider the set of paths in $\Gamma(W, S)$ of length n which start at the identity. The probability $P^n(w)$ that the walk is at w after n steps is equal to the proportion of these paths which end at w . Lazy walks can be considered by including 1_W in S .

2.2 Distances in $\Gamma(W, S)$

It will be important to understand relative distances in the Cayley graphs of Coxeter groups. Consider the Cayley graph with the usual graph metric — that is, each edge has length 1, and the distance $d(w, x)$ between two vertices w and x is the number of edges in the shortest path connecting them. As usual in the theory of Coxeter groups, $l(w)$ will denote the distance from the identity $d(1_W, w)$. Equivalently, $l(w)$ is the fewest number of generators which can be multiplied to produce w .

For this section, fix $w \in W$ and $s \in S$, with $l(w) < l(ws)$. That $l(w) < l(ws)$ is not used in this section, but is consistent with how these results will be used in Section 3.

Definition 6. Let $\Gamma(W, S)$ be the Cayley graph of a Coxeter system (W, S) . For the fixed adjacent vertices w and ws of $\Gamma(W, S)$, colour each vertex of $\Gamma(W, S)$ white if it is closer to w than to ws and black if it is closer to ws than to w .

Proposition 7. Each vertex of $\Gamma(W, S)$ is coloured white or black, but not both.

Proof. Each vertex of $\Gamma(W, S)$ has at most one colour, because it cannot be both closer to w than to ws and the reverse. To show that each vertex is coloured, it is sufficient to show that no vertex can be equidistant from w and ws .

The Coxeter relations of (W, S) each have even length, so $\Gamma(W, S)$ is a bipartite graph, and hence the distances from any vertex to w and ws have opposite parities. Thus each vertex is coloured, completing the proof. \square

Definition 8. Continuing from Definition 6, colour grey each edge which connects a white vertex to a black vertex.

Lemma 9. If (x, xt) is a grey edge, with x white and xt black, then $d(x, w) = d(xt, ws)$. (The generator t may be equal to s , but need not be.)

Proof. The two vertices w and ws are adjacent, as are the vertices x and xt . The vertex x is white, and xt is black. Thus, the following relations between distances hold

$$\begin{aligned} d(x, w) + 1 &= d(x, ws) \\ d(xt, ws) + 1 &= d(xt, w) \\ d(x, ws) &= d(xt, ws) \pm 1 \\ d(xt, w) &= d(x, w) \pm 1 \end{aligned}$$

Adding these four equations, each of the two \pm signs must be a plus. Thus $d(x, w) = d(xt, ws)$, as required. \square

Lemma 10. *Under the conditions of Lemma 9, $w^{-1}x = sw^{-1}xt$.*

Proof. Let ω be a reduced word for $w^{-1}x$. From Lemma 9, $l(sw^{-1}xt) = l(w^{-1}x)$. Thus, $s\omega t$ is a word of length two greater than the minimum length of any equivalent word. By the deletion condition (Section 1.7 of [7]), there is a reduced word for $sw^{-1}xt$ which can be obtained by deleting two letters from $s\omega t$. However, Lemma 9 also implies that the words $s\omega$ and ωt are reduced, so the two letters removed from the word $s\omega t$ must be the initial s and the final t . Therefore $s\omega t$ and ω are equivalent words, so $w^{-1}x = sw^{-1}xt$. \square

Proposition 11. *The map $L_{ws w^{-1}}$ interchanges the endpoints of any grey edge.*

Proof. Let (x, xt) be a grey edge, with x white and xt black. The image $L_{ws w^{-1}}(x)$ is

$$\begin{aligned} L_{ws w^{-1}}(x) &= wsw^{-1}x \\ &= wssw^{-1}xt \text{ (by Lemma 10)} \\ &= xt \end{aligned}$$

The map $L_{ws w^{-1}}$ is an involution, so $L_{ws w^{-1}}(xt) = x$. \square

Corollary 12. *Any vertex of $\Gamma(W, S)$ is adjacent to at most one grey edge.*

Proof. The function $L_{ws w^{-1}}$ is well defined. If any vertex were adjacent to more than one grey edge, then Proposition 11 implies that $L_{ws w^{-1}}$ is multivalued, a contradiction. \square

The results in this section are not new — they are standard facts about the geometry of $\Gamma(W, S)$, proven again here to draw out the key pieces. The set of grey edges is commonly referred to as a wall, and could be defined as the set of edges preserved by the reflection $L_{ws w^{-1}}$.

2.3 The weak Bruhat order

Let (W, S) be a Coxeter system. The (right) weak Bruhat order on (W, S) is defined as follows (Chapter 3 of [4]).

Definition 13. *Let w and w' be elements of W . Then $w \leq_B w'$ if there is a reduced word for w which can be multiplied on the right by elements of S to produce a reduced word for w' .*

An equivalent formulation in terms of the Cayley graph of (W, S) is

Definition 14. *Let w and w' be elements of W . Then $w \leq_B w'$ if there is a minimal length path in $\Gamma(W, S)$ from the identity element 1_W to w' which passes through w .*

These two definitions are equivalent because edges in the Cayley graph $\Gamma(W, S)$ correspond to multiplication on the right by an element of S .

3 Proof of main result

The main result of this paper is that the weak Bruhat order arises as a likelihood order for random walks on (W, S) .

Theorem 1. *For any Coxeter system (W, S) , consider the simple random walk on W , starting at the identity, and at each step multiplying on the right by an element of S or by the identity, each with probability $\frac{1}{|S|+1}$. Then for any n , and any two states w and w' , if $w \leq_B w'$ then the probability that the random walk is at w after n steps is at least the probability that it is at w' .*

To prove Theorem 1, it suffices to consider w and w' which are adjacent — that is, when $w' = ws$ for some $s \in S$, with $l(w) < l(ws)$. If the result is true for adjacent vertices, then the general case follows by induction. Thus, the theorem reduces to the following proposition.

Proposition 15. *Let $w \in W$ and $s \in S$, with $l(ws) > l(w)$. Then for any n , $P^n(ws) < P^n(w)$.*

Proof. Let \mathcal{P} be the set of paths of length t from the identity 1_W to w , with each step being either the identity or an element of S . Let \mathcal{P}' be the set of paths of length t from 1_W to ws . To prove this proposition, it suffices to construct an injection from \mathcal{P}' to \mathcal{P} . Consider an element α of \mathcal{P}' , and write it as

$$\alpha = (1_W = a_0, a_1, a_2, \dots, a_n = ws).$$

The path α starts at a point closer to w than to ws , because $l(ws) > l(w)$, and α ends at ws , which is closer to ws than to w . Thus at some point, α crosses from a vertex closer to w to a vertex closer to ws — that is, this path crosses a grey edge. Given α , let i be the last time at which α either just crossed a grey edge or just stayed in place on an endpoint of a grey edge. Using \oplus to denote concatenation, define the sequence of vertices

$$\begin{aligned} f(\alpha) &= (a_j)_{j=0}^{i-1} \oplus (L_{ws w^{-1}}(a_j))_{j=i}^n \\ &= (1_W = a_0, a_1, \dots, a_{i-1}, L_{ws w^{-1}}(a_i), L_{ws w^{-1}}(a_{i+1}), \dots, L_{ws w^{-1}}(a_n)) = w. \end{aligned}$$

That is, the sequence $f(\alpha)$ is identical to α up until time $i - 1$, and from time i onwards, it is reflected by $L_{ws w^{-1}}$.

Proposition 16. *The sequence $f(\alpha)$ is a path. That is, each two consecutive entries are either adjacent in the graph $\Gamma(W, S)$, or equal.*

Proof. It must be checked that each two consecutive entries in $f(\alpha)$ are either adjacent in the graph $\Gamma(W, S)$, or equal. This is immediate for each consecutive pair except for a_{i-1} and $L_{ws w^{-1}}(a_i)$.

By the definition of i , the vertex a_i is a black vertex adjacent to exactly one grey edge, and a_{i-1} is one of the two endpoints of that edge. From Proposition 11, $L_{ws w^{-1}}(a_i)$ is either equal to a_{i-1} or connected to a_{i-1} by this grey edge. \square

The map $L_{ws w^{-1}}$ interchanges ws and w , so f is a function from \mathcal{P}' to \mathcal{P} . All that remains is to show that f is an injection. The function f is an involution, because it applies $L_{ws w^{-1}}$ to the part of the path from time i onwards, the map $L_{ws w^{-1}}$ is

an involution, and moving from α to $f(\alpha)$ does not change the definition of i (but rather interchanges the two cases in the definition of i).

Thus f is an involution from \mathcal{P}' to \mathcal{P} , so there are at least as many paths of length n from 1_W to w as from 1_W to ws , completing the proof of Proposition 15. \square

Corollary 17. *For any Coxeter system (W, S) and the corresponding random walk described by Theorem 1, the most likely element after n steps is the identity. If W is finite, then the least likely element is the longest element. Here, most and least likely are not necessarily strict.*

Example 18. Consider the random walk on S_4 generated by the adjacent transpositions $(1\ 2)$, $(2\ 3)$, and $(3\ 4)$, as well as the identity. For any n , the most likely element after n steps is the identity and the least likely is the reversal $(1\ 4)(2\ 3)$. Theorem 1 does not address the relative likelihoods of the transpositions $s = (1\ 2)$ and $t = (2\ 3)$, but both are more likely than $(1\ 3) = sts = tst$.

Example 19. Take the simple random walk on the square grid $\mathbb{Z} \times \mathbb{Z}$ generated by $(\pm 1, 0)$, $(0, \pm 1)$ and $(0, 0)$. This is a relabelling of the Cayley graph of the Coxeter group

$$D_\infty \times D_\infty = \langle s, t, u, v \mid s^2, t^2, u^2, v^2, susu, svsv, tutu, tvtv \rangle.$$

For any n , the most likely element after n steps is the identity $(0, 0)$, and the next most likely are the four adjacent vertices (which are equally likely, by symmetry). The vertex $(1, 3)$ is always more likely than $(2, 3)$, but Theorem 1 doesn't address the relative likelihoods of $(1, 3)$ and $(2, 2)$.

Example 20. Finally, consider the simple random walk on a d -regular tree, with laziness $\frac{1}{d+1}$. This is the Cayley graph of the free product of several copies of the group $\mathbb{Z}/2\mathbb{Z}$. For any n , the most likely vertex after n steps is the initial vertex, the next most likely are the adjacent vertices, then the vertices at distance two, and so on.

Theorem 1 may be strengthened to not require that all generators have equal probability.

Theorem 21. *For any Coxeter system (W, S) , consider a random walk on W , starting at the identity, and at each step multiplying on the right by an element $s \in S$ or by the identity, with probabilities p_s or p_{id} . As long as each p_s is less than p_{id} , the conclusion of Theorem 1 holds.*

Proof. The only part of the proof that must be changed is the proof of Proposition 15, comparing the probabilities of the states ws and w , for $w \in W$ and $s \in S$ with $l(ws) > l(w)$. With this choice of s , divide the probability of multiplying by 1_W into two parts, of probabilities p_s and $p_{\text{id}} - p_s$. Where the proof of Proposition 15 pairs up the events of multiplying by s or by 1_W , use only the first of these parts, which is an event of equal probability to that of multiplication by s . \square

References

- [1] Sami Assaf, Persi Diaconis, and K. Soundararajan. A rule of thumb for riffle shuffling. *Ann. Appl. Probab.*, 21(3):843–875, 06 2011.
- [2] M. Bernstein. Likelihood Orders for some Random Walks on the Symmetric Group. *ArXiv e-prints*, June 2013, 1306.5008.
- [3] M. Bernstein. A random walk on the symmetric group generated by random involutions. *ArXiv e-prints*, June 2016, 1606.09588.
- [4] A. Bjorner and F. Brenti. *Combinatorics of Coxeter Groups*. Graduate Texts in Mathematics. Springer Berlin Heidelberg, 2006.
- [5] P. Diaconis and M. Isaacs. Least likely elements for random walk on finite groups. *Unpublished*.
- [6] Persi Diaconis and Laurent Saloff-Coste. Comparison techniques for random walk on finite groups. *Ann. Probab.*, 21(4):2131–2156, 10 1993.
- [7] J.E. Humphreys. *Reflection Groups and Coxeter Groups*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1992.
- [8] Johan Jonasson. Mixing times for the interchange process. *Latin American Journal of Probability and Mathematical Statistics*, IX:667–683, 2012.